

Toward a Computational Model of Narrative

George Lakoff, Sridhar Narayanan

International Computer Science Institute and
University of California at Berkeley
lakoff@berkeley.edu
snarayan@icsi.berkeley.edu

Abstract

Narratives structure our understanding of the world and of ourselves. They exploit the shared cognitive structures of human motivations, goals, actions, events, and outcomes. We report on a computational model that is motivated by results in neural computation and captures fine-grained, context sensitive information about human goals, processes, actions, policies, and outcomes. We describe the use of the model in the context of a pilot system that is able to interpret simple stories and narrative fragments in the domain of international politics and economics. We identify problems with the pilot system and outline extensions required to incorporate several crucial dimensions of narrative structure.

Introduction

We structure our lives by narratives and we understand events in the world in terms of narratives — events of all kinds, in science, in politics, in every facet of life. Some of these are conscious, others are very much unconscious. In life narratives, each of us is the protagonist, living out the narratives as best we can.

Computer systems that attempt to model narrative need to access the underlying cognitive structure of human motivation, actions, goals, and events. To illustrate the scope of the challenge, consider the following simple narrative from a recent article (May 2010) in the online version of the Wall Street Journal about the state of the US economy.

The US Economy is on the verge of stumbling back into recession. The jobs picture is dismal, and the one-time boost from the stimulus package is almost over. The stimulus could have been better spent to buttress the economy through job growth.

For humans, the passage above is easy to understand, to reason about, make decisions on, and take appropriate actions (such as not trying to change jobs). Current computer systems, on the other hand, can only reason with precise semantics, and currently have no way to represent or interpret phrases such as “on the verge of”, “stumbling back”, “dismal”, “boost”, “almost over”, “buttress”, “could have been”. Let us look at a few of the semantic distinctions

made to get a better appreciation of the structural, semantic, and qualitative nature of language. First, note that institutions (US) are conceptualized as causal agents, abstract actions (economic state changes) as physical motions (stumbling), causation or state maintenance as forces (but-tress). Also abstract states (recession) are expressed as features (holes) in a spatial terrain. These mappings are part of a larger composite metaphorical mapping, called the Event Structure Metaphor [7,11,12], that projects inferences from physical motion and manipulation to abstract actions, goals, and policies and is common in all languages studied to date.

Second, notice the complex, dynamic, fine-grained, and context sensitive scenario information contained in the phrase “on the verge of stumbling back”. “On the verge” suggests that that the event (stumbling) has not yet started but is likely to start soon. The word stumble here encodes a fairly complex scenario where an ongoing economic policy has encountered some difficulty. The VP+ing (stumble +ing) construction denotes an ongoing situation, which could likely lead to a failure of the policy (just as physical stumbling could lead to falling). Also the phrase “back” indicates that the economy was in a similar state (recession) in the recent past. As a whole, the phrase “on the verge of stumbling back” suggests that if the current state of affairs continues (both the policy and the economic environment), then the economy is likely to have negative growth as it did before the current recovery. The reader is invited to consider the metaphoric meaning of “dismal” in the context of the sentence.

Third, language routinely involves imagining alternative goals, resources, outcomes, and actions (or inactions). Consider the counterfactual implication of the subjunctive “could have been”. Counterfactuals are mental simulations of “variations on a theme”. They refer to imagined alternatives to something that has actually occurred. Counterfactual reasoning is basic to human cognition and is ubiquitous in commonsense reasoning as well as in formalized discourse. They play a significant role in other cognitive processes such as conceptual learning, planning, decision making, social cognition, mood adjustment, and performance improvement. In previous work [14,19], we present a modeling framework and results that represent the first step

toward a computationally adequate cognitive model of counterfactuals. Our treatment of counterfactuals comes from independent considerations of cognitively motivated event structure representation useful for event coordination and for language processing. Our model is able to capture the variety, scope, and inferential richness of the psychological data and makes detailed predictions about the neural substrate that underlies counterfactual processing.

More generally, reliance on fine-grained, context sensitive information about human goals, actions, policies, and outcomes is central to the semantics of language. Computer systems cannot make headway in processing human narrative unless they can represent and reason with such information. We address this challenge by proposing computer systems that simulate human knowledge and inference. Our approach combines over two decades of interdisciplinary work within the Neural Theory of Language (NTL) (<http://www.icsi.berkeley.edu/NTL>) and FrameNet (<http://framenet.icsi.berkeley.edu>) projects at the International Computer Science Institute (ICSI) and the University of California, Berkeley. The effort combines results from cognitive linguistics and computer science to build systems capable of meaning. This paper focuses on the structure and interpretation of narrative, describes the current state of our results, and identifies key directions of future work on the multiple dimensions of narrative.

Components of a cognitive model of narrative

Proposition: *Narrative exploits the rich structure of human event and action representation. Encoding this structure is necessary for representing, reasoning about the form and content of narratives.*

A general ontology capable of describing human actions and events must fulfill some essential requirements. The action ontology and corresponding model has to be a) fine-grained to capture the wide range of possible events and their interactions; b) context-sensitive and evidential in order to adapt to a dynamic and uncertain environment; c) cognitively motivated to allow humans to easily query and make sense of the answers returned; and d) elaboration-tolerant so that new domain models can specialize existing representations without changing the basic primitives.

We have developed a parameterized model of the structure of events and processes that meets these requirements. Figure 1 shows the basic schema of events. We describe its main elements here.

All events have a basic structure: A basic event is comprised of a set of inputs, outputs, preconditions, effects (direct and indirect), and a set of resource requirements (consuming, producing, sharing and locking). Events are grounded at a time and place and have a duration. The *hasParameter* link in Figure 1 depicts the set of parameters

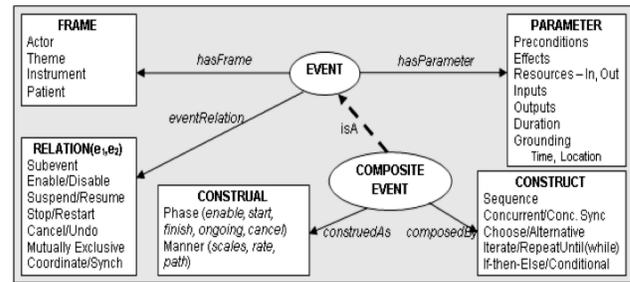


Figure 1: An ontological schema of Event Structure.

in the domain of the basic event type.

Composite events have rich temporal structure and evolution trajectories: The fine-structure of events is composed of key states (such as enabled, ready, ongoing, done, suspended, canceled, and stopped) and a partially ordered directed graph of transitions that represents possible evolution trajectories between these states (transitions include: prepare, start, interrupt, finish, cancel, iterate, resume, re-start). Each of these transitions may be atomic, timed, stochastic or hierarchical (with a recursively embedded event-structure).

Composite events are composed of process primitives: Verbal aspect (the temporal structure of events [11]) discriminates between events that are punctual, durative, (a)telic, (a)periodic, (un)controllable, (ir)reversible, ballistic, or continuous. Each type of event relates to a particular internal structure, which draw upon a set of process primitives and control constructs (sequence, concurrent, choice, conditionals, etc.). These primitives specify a partial execution ordering over subevents. The *composedBy* relation in Figure 1 shows the various process decompositions. [11,14,19] describe them in greater detail.

Composite events support various construals: Composite events can be viewed at different granularities using operations for *elaboration* (zoom-in) and *collapse* (zoom-out) [11]. In addition, specific parts and participants of a composite event can be *focused on*, *profiled* and *framed*. Construal operations are shown in Figure 1 through the *construedAs* relation.

Events relate to each other in regular patterns: A rich theory of inter-event relations allows sequential and concurrent enabling, disabling, or modifying relations. Examples include interrupting, starting, resuming, canceling, aborting or terminating relations, as shown in Figure 1 through the *eventRelation* relation.

Modeling Narrative: A Pilot System

Complex reasoning about event interactions requires not only an event description, but also a dynamic model that can simulate the execution of the event unfolding over time. We can instantiate such a model with facts about a particular event, enabling us to project which situations are likely or possible based on the consumption and production of resources and the creation and elimination of states.

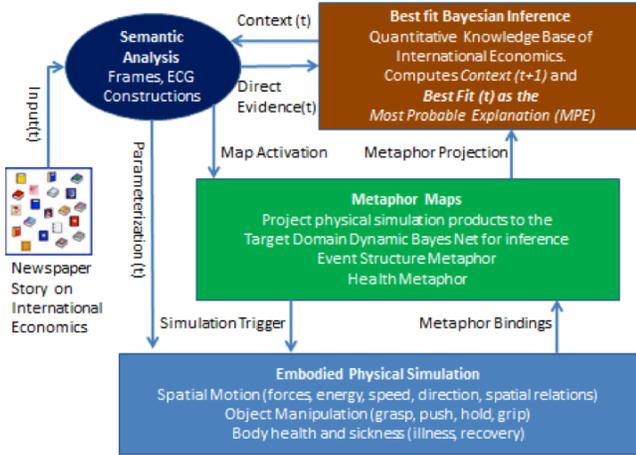


Figure 2: The KARMA pilot system for story understanding. The system interprets newspaper story fragments in the domain of international economics in context using metaphorical projections of embodied simulations from the domains of spatial motion and manipulation onto the abstract domain of economic policies, goals, and outcomes. The system computes the probabilistic “best-fit” of the input utterance, in the context of a) the background knowledge of the target domain (economics), b) previous utterances, c) metaphorical projections, and d) the evolving situation. The pilot system did not have the semantic analysis component which was more recently implemented [1,3,6] and not yet integrated.

Figurative reasoning about narratives requires modeling coordinated temporal processes complex, structured mental states. Our action theory for the Pilot system comprised of two central components; 1) an *active* (executing) representation of actions and events (called X-nets) based on extensions to Petri Nets and 2) a Temporal (Dynamic) Bayes Net model of *state* that captures and reasons about complex dependencies between state variables.

Systematic metaphors project these features onto abstract domains such as Economics enabling language to use motion terms to describe abstract actions and processes. The implemented system has three main components, namely the *source domain*, the *target domain* and the *metaphor maps*. The source and target domains are based on a model of action that is able to meet the representational requirements and support the kinds of inferential processes inherent in language understanding.

The central idea behind the model is that the reader interpreting a phrase that corresponds to a motion term is in fact performing a mental simulation of the entailed event in the current context. The basic idea is simple. We assume that people can execute simulations with respect to structures that are not linked to the body, the here and the now. In this case, actions are not carried out directly, but instead trigger simulations of what they would do in the imagined situation. The physical world is modeled by other schemas that have I/O links to the schema representing the planned action.

In the pilot implementation, source domain structure is encoded as connected motion schemas. The model of the source domain is a dynamic system based on inter-X-net

activation, inhibition and modulation. In the simulation framework, whenever an executing x-net makes a control transition, it potentially modifies state, leading to asynchronous and parallel triggering or inhibition of other x-nets. The notion of state as a graph marking is inherently distributed over the network, so the working memory of an x-net-based inference system is distributed over the entire set of x-net. This control and simulation regime remains central to the proposed CPRM design. Of course, it is also intended to model the massively parallel computation of the brain [5].

An important and novel aspect of our source domain representation is that the same system is able to respond to either direct sensory-motor input *or* other ways of setting the agent state (such as linguistic input). This allows for the same mechanism to act and perform inference through imaginative simulation. The same X-net circuit can also be used for high-level control and reactive planning. There is now robust biological evidence to support the view [7] that planning, recognition and imagination share a common representational substrate. Our computational model, which we call *simulation semantics* [14] is largely motivated by constraints of neural computation [5] and is offers a computational substrate for these findings. We believe this to be an important aspect of embodiment allowing the same mechanisms to reason as well as to act.

The structure of the pilot abstract domain (the domain of international economic policies) encodes knowledge about economic policies. The representation must be capable of a) representing background knowledge (such as the US is a market economy), b) modeling inherent target domain structure and constraints (high-growth may result in higher inflation), and c) be capable of computing the impact of new observations which may from direct input (“US economy is experiencing high-growth”), or from metaphoric (or other) inferences (“economy stumbling”). Furthermore, these different sources of evidence have different degrees of believability, and the representation must provide a framework for their combination.

A story represents the specification of a *partial* trajectory over epistemic states modeled by the Dynamic Bayes Net. This is simulated by clamping some of the Bayes network nodes to specific values. The remaining features are estimated using known target domain inter-feature correlations as well as metaphoric projections from the embodied general knowledge (x-nets). Metaphoric projections of x-net executions may clamp target features to specific values conditioning (by placing new evidence on) the target domain Bayes net.

Comprehending a story corresponds to finding the set of trajectories that best satisfy the constraints of the story and are consistent with the domain knowledge. This may involve *filling in* missing values or placing new evidence on the Bayes net. The resultant target network state becomes a prior for processing the next input at stage $t = 2$. Background knowledge is encoded as the network state at $t = 0$.

Target inferences can go forward and backward in time in the estimation of the *best fitting interpretation* (most probable explanation) of the input story.

In the pilot system, the embodied domain theory had about 100 linked x-nets, while the abstract domain theory is a relatively sparse net of about 40 multi-valued variables with at most 4 temporal stages. It also encoded about 50 metaphor maps from the domains of *health* and *spatial motion*. These were developed using a database of 50 2-3 phrase fragments from newspaper stories all of which have been successfully interpreted by the program. Among the inferences made were those related to goals (their accomplishment, modification, subsumption, concordance, or thwarting), resources, aspect, frame-based inferences, perspectival inferences, and inferences about communicative intent. [3,5,19] report on the different types of inferences produced by the system.

In summary, our results suggest that a large proportion of commonplace narratives of abstract events seem to project embodied, familiar concepts onto more abstract domains such as economics and politics. This allows non-experts to comprehend and reason about such abstract policies and actions in terms of more familiar and universal experience. The fact that the metaphoric inferences are context-sensitive, immediate, and defeasible set up fairly strong representational requirements for a metaphor interpretation system. The structured probabilistic representation coupled with the rich action semantics of x-net based simulation enables our model to capture subtle contingency relations between events necessary for routine common-sense inference. To make the system and its results available to the broader community, some fundamental problems need to be addressed.

Problems with the Pilot System

The KARMA narrative interpretation system was a “proof of principle” demonstration. While the variety and subtlety of inferences made exceeded any other system we are aware of, there was not a detailed computational analysis of the scale, scope and quality of the information communicated through metaphoric language or the ability of our approach to perform these inferences. One major barrier to scaling the pilot system to more complex inference tasks is the inefficiency of very large unstructured Bayes networks. A central problem with our DBN based state representation is that it is propositional and does not scale well to relational domains. We have been addressing this issue with a relational state representation based on Probabilistic Relational Models (PRM). Our approach, called Coordinated Probabilistic Relational Models (CPRM) is a synthesis of Stochastic Petri Nets for action and Probabilistic Relational Models for inference. Leon Barrett’s thesis [2] describes CPRM, and its use in real time action modeling and inference. A full description of the computational architecture of CPRM is outside the scope of this paper.

A second problem is directly related to this paper and deals with the inadequate treatment of narrative structure within the pilot system. While narratives make use of event structure and goals and actions, narratives have specific dimensions and structure that the pilot system does not address. The point is that a narrative is far more than just a description of events. A narrative has a cognitive structure, and a given narrative may extend over time, often a long time. Newspaper reports are often about stages in a narrative as was shown in the pilot system results.

Narratives can structure the past (as in autobiography and explanation of the current situation) or the future (what to expect and what to do), and in the present they can link the past to expectations about the future. Narratives tell you what is important and why. A narrative makes certain experiences salient, that is, it gives them a value, by activating them in a structure. The salient structure can be a complication (say, a threat), a denouement (say, a call for action), or a resolution (say, an occasion for satisfaction). Narratives are called on to satisfy curiosity, allow for empathy and self-projection, guide memory formation and reconstruction, and provide an autobiographical self. Narratives allow us to function sensibly in the world and are central to a sense of self. It is time to look at precisely what structures make up a narrative and how we can model them and study them scientifically.

Dimensions of Narrative Structure

The scientific study of narrative requires five aspects of narrative structure.

- The dimensions of structure in elementary narratives.
- The compositional principles governing how elementary narratives combine to form complex narratives.
- The conceptual metaphors that map basic narrative structures onto many subject matters — from fairy tales, to detective stories, to politics, to stories of scientific discovery.
- The principles of linguistic pragmatics governing how complex narratives are told in context.

Let us begin with the dimensions of structure in elementary narratives. We will then discuss some compositional principles for complex narratives.

Dimensions of structure in Narratives

Dimension 1: Moral systems and guides to living

Fables and stories typically have morals. We see these overtly by the dozen in Aesop’s fables. Their morals are ways of understanding the world and guides to both moral and practical living. Consider the fable of “The Bat, the Birds, and the Beasts.”

A great conflict was about to come off between the Birds and the Beasts. When the two armies were collected together the Bat hesitated which to join. The Birds that passed his perch said: "Come with us"; but he said: "I am a Beast."

Later on, some Beasts who were passing underneath him looked up and said: "Come with us"; but he said: "I am a Bird." Luckily at the last moment peace was made, and no battle took place, so the Bat came to the Birds and wished to join in the rejoicings, but they all turned against him and he had to fly away. He then went to the Beasts, but soon had to beat a retreat, or else they would have torn him to pieces. "Ah," said the Bat, "I see now, "He that is neither one thing nor the other has no friends."

Aesop's fables are all metaphorical, nominally about animals, but really about people. The fundamental metaphor is that *Human Characteristics Are Animal Instincts*, and the point of the fables is that humans act like animals, but with insight, humans can make choices to change their circumstances, whereas animals cannot.

Narratives typically have a moral dimension. Moral systems structure systems of narratives, which allow narratives to provide guidelines for how to live. Villains must be punished, heroes rewarded, the Horatio Alger hero should succeed, the tragedy results in harm to the protagonist, and so on. Conceptual metaphor is central to the constitution of moral systems and to projection of narratives onto everyday situations. The KARMA system, as it stands, needs to be enhanced with moral narrative structures and conceptual metaphors pertaining to morality.

Dimension 2: Folk Theories of how people and things work.

Folk theories are largely unconscious, automatic cognitive structures characterizing how things work or what properties things and people have. There is nothing derogatory about our use of the term "folk." They are used in everyday life. There are folk theories about what people are like, what causes what, what are plans and goals, what and why people steal, why people buy, how light switches work, what intelligence is, how people learn, how politics works, and so on. Part of what constitutes a culture or subculture is its collection of folk theories and the logics they bring with them into narratives. Folk theories provide a crucial backdrop to narrative structure, and they are used to draw morals and other inferences. For example, take folk theories of learning. Here are some common ones: (A) People learn only when rewarded for learning and punished for not learning. (B) People are naturally inquisitive and learn on their own when obstacles are removed. (C) People learn when they have good teachers. (D) People learn when good theories of learning are applied. All of these show up in one type of narrative or another. We believe that the rich representation of events, goals, outcomes, and behavior underlying the KARMA system architecture provides a promising framework to encode folk theories and their use in narrative; however the knowledge engineering task remains an ongoing endeavor.

Dimension 3: Overall Plot Structure.

Narratives have a high-level organizational structure that, in typical cases, looks like this: Generic plot roles (e.g., Protagonist, Antagonist, Helpers), a Background, a

Complication, a Main Event (e.g., a struggle, test, trial, decision, or other crucial event), a Denouement (that is, a Resolution), the Consequence, and the Moral (if any). Conceptually, these are linearly ordered in the conceptual logic of the plot.

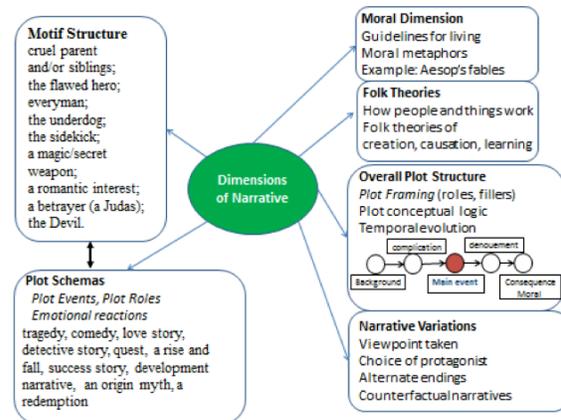


Figure 3: The dimensions of narrative structure.

In an actual story told in language, they may be ordered differently. For example, a newspaper story might lead with any one of these, depending on what constitutes "news." For example, The NY Times on July 31, 2010, included the following stories exemplifying different parts of overall plot structure.

- "Voice on Phone Is Lifeline for Suicidal Veterans" is a *Background* story about problems of veterans fitting back into society after serving in the Middle East.
- "Debate Heating Up on Plans For Mosque Near Ground Zero" is about a *Complication*: The Anti-Defamation League has decided to oppose the mosque.
- "Flu Vaccines Are Approved and Urged for Most" reports on a *Main Event*, a decision by the Food and Drug Administration to approve flu vaccines and recommend them to the public.
- "Advice By Panel Is To Reprimand, Not Oust, Rangel" is a *Denouement* to the process of investigating charges against Rep. Charles Rangel, but it is a *Complication* for the next step in the process.
- "Afghan Women Fear the Loss of Modest Gains" is a *Consequence* of the Taliban resurgence in Afghanistan.

KARMA currently does not encode plot structures and would have to be extended with a plot structure X-net that captures the conceptual logic of the dynamically unfolding narrative and generates expectations and inferences in a manner very similar to the *controller* network in our model of linguistic aspect [13,14].

Dimension 4: Plot Schemas

Specific plots fill in the overall plot structure of a narrative with more specific content. Examples include:

- Tragedy, comedy, hero story, love story, detective story, quest, a rise and fall, a success story, a development narrative, an origin myth, a redemption, and so on.

- Each specific plot has specific events that fill in the parts of the overall plot structure; it has to have a Background (*Once upon a time, In the Colorado primary, and so on*), a Complication (A popular legislator challenges the wealthy incumbent senator who has raised more money), Main Event (The primary election), etc.
- Each specific plot has Plot Roles, e.g. Hero, Villain, Victim; Hero on a quest, his or her Holy Grail, Difficulties along the way; and so on.
- Conventional emotional reactions to each specific plot event; e.g., anger at villainous action, anxiety and fear at encounter of hero with villain; relief and joy at hero's victory; satisfaction with the consequence.

Dimension 5: Motif structure

Narratives are often used to understand and describe one's life. This idea has been a commonplace in psychotherapy since Freud's description of the Oedipus Complex. A recent example is Collette Dowling's book, *The Cinderella Complex*. There she describes a metaphor based on the structure of the Cinderella fairy tale. Cinderella is beautiful, good-hearted, bright, and hard-working, but she is oppressed by the situation she is in. She is a helpless stepchild oppressed by an evil stepmother and two evil stepsisters. She cannot escape her situation on her own. She has to be saved by a man — the Prince. In the Cinderella narrative that Dowling describes, attractive, good-hearted, talented women see themselves as Cinderellas, taking themselves to be helplessly caught in oppressive situations, waiting for their prince to come. They are living their lives by the Cinderella narrative.

What we learn from this is that classical narratives can be applied metaphorically to one's own life or to other situations, often unconsciously. Another thing we learn is that there are what the great folklore scholar Stith Thompson called "motifs." Our culture has a great many. The Cinderella figure is only one. Another is the Devil, embodying pure evil, out to lure moral people into doing immoral things at the cost of their immortal souls. In literature, there is Faust and The Devil and Daniel Webster. In real life, the devil motif, a villain of pure evil, has been used of Charles Manson and Saddam Hussein.

A favorite motif is the hero with a fatal flaw that threatens to, or does, lead him to a tragic end. Achilles is the classical case: the strongest, bravest, best-looking Greek warrior, whose body is invulnerable except for his heel. He wins battle after battle, until in a crucial battle he is shot in the heel with an arrow and dies. A modern version is Superman, who is vulnerable only to Kryptonite, which villains somehow get a hold of. The fatal flaw usually appears in the course of a heroic quest — whether to defeat the Trojans and bring Helen back to Greece or to rid Metropolis of criminals and occasionally save the world.

Well-known conventional types filling in plot roles, described extensively in Stith Thompson's *Motif Index* (<http://www.folklore.bc.ca/Motifindex.htm>), include the cruel parent and/or siblings; the flawed hero, everyman, the

underdog, the sidekick; a magic or secret weapon; a romantic interest; a betrayer (a Judas); the Devil; and so on. Such choices often come with even more specific plot schemas. These figures usually come with more specific plot structures. The Devil tries to get someone's immortal soul by offering to fulfill certain desires. The Cruel Parent and Siblings often goes with a virtuous, but oppressed Cinderella who is saved by a Prince who recognizes her beauty and virtue. The Underdog is a virtuous Hero who has to compete with a much more powerful Villain.

In future work, we plan to use the FrameNet and ECG structure [4,5,6] to encode the commonly occurring motifs and the X-net computational framework to encode the different plot schemas. Our previous work [4] combining these techniques for frame based inference in KARMA points toward a possible integration of motif structure and plot schemas in a computational model of narrative.

Dimension 6: Narrative Variations

Narratives may vary the viewpoint taken, the choice of protagonist (who you identify with), choose an alternative ending to a classic narrative, and so on. For example, John Gardner's *Grendel* tells the story of Beowulf from the point of view of the monster, Grendel. In *The Yiddish Policemen's Union*, after the Holocaust, Israel doesn't work out in 1948 and a large number of Jews are relocated for 60 years to Sitka, Alaska. A Hasidic sect is a violent gang. And the potential Messiah decides he doesn't want to be the Messiah. The narrative is structured by variations on more traditional narratives. The Israel Bond novels are comic take-offs on James Bond novels, with the hero being a Jewish schlep named Israel Bond. Recent extensions to the KARMA system with algorithms for inference with alternative and counterfactual narrative [14] have already shown how some of the variations can be modeled. A proper treatment of narrative viewpoint will require the full integration with mental spaces and ECG grammar [1,3,5,6], which is an ongoing task.

Plot Composition

Complex narratives are often composed of elementary narratives. These were first described in George Lakoff's 1964 [9] address to the Linguistic Society of America, "Structural Complexity in Fairy Tales." It follows up on Vladimir Propp's classic "Morphology of the Folktale" [15], reanalyzing the data there in terms of Chomskyan generative grammar. This was followed up on in the 1970's by David Rumelhart's work on story grammars and Schank and Abelson's scripts [16,17,18].

Figure 4 shows some basic plot composition structures. The composition structures are a subset of the existing composition templates in KARMA and its extensions

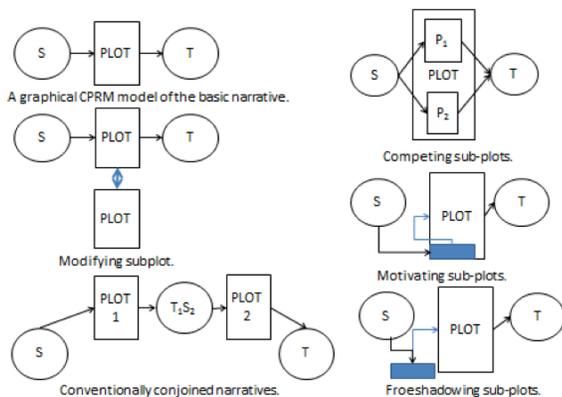


Figure 4: Some examples of compositional narrative structures. The basic narrative is modeled as a CPRM module (please see the text) with a narrative start (S) and end (T) state. The Plot CPRM (shown as a rectangle) encodes the various participants and roles and simulates the evolving narrative situations and events. The various arrangements of narratives, sub-narratives, and plots/sub-plots constitute the various compositional possibilities (see [13,14] for more models of composition).

[2,12,13,14] and are already available for narrative modeling. The basic narrative is modeled as a CPRM and the different compositional principles specify control and data arrangement of individual plot and narrative schemas. The most typical compositional principles are:

- *Conventionally conjoined subnarratives*, e.g. the hero goes first to the Copper Kingdom, then the Silver Kingdom, then the Gold Kingdom; the candidate first loses a minor race, learns his lesson, and goes on to win successive races; someone on a quest must overcome a succession of obstacles before getting a chance at the Holy Grail; and so on. Each sub-narrative is a predecessor in a sequence of narratives.
- *Modifying subplots*. For example, in Russian fairy tales, there is often a subplot that tells how the hero, Ivan, got his magic sword, a subplot in which he encounters the Baba Jaga and answers three riddles and gets the sword as a reward. In the narrative of the Vietnam War, the story of the publication of the Pentagon Papers is a modifying subplot, part of the story of how the public turned against the war. Modifying subplots may play a role in a larger plot or be used to exemplify the qualities of the hero, villain, or victim.
- *Competing subplots*. A competition may be understood through competing elementary narratives, where the hero of one is the villain of the other. The overall narrative may have one as the conquering hero, with the two subplots converging at the Main Event, the competition of the two heroes.
- *Motivating subplots*: A subplot may motivate the main plot, say, by reporting on some villainy with the main plot reporting on the revenge for it, or in general giving a rationale for the main plot.
- *Foreshadowing subplots*: A preceding subplot may have the same structure as the main plot, but with

minor characters. Or the subplot may exemplify the threat to the hero in the plot.

This is not intended to be a complete list. But it should provide an idea of how complex narratives arise from elementary narratives.

Metaphor projects narrative structure

Conventional stories with conventional motifs are often mapped by conceptual metaphors onto everyday situations, e.g., in politics, science, the arts, a personal biography. The KARMA system discussed earlier is an implemented computational model that captures the essential components of this phenomenon. In addition, [12] hypothesizes specific invariants that project aspects of familiar and embodied human experience to structure stories about economics, politics, science and more abstract domains. These hypothesized invariant experiential structures are called *Cogs*. Cogs are candidate neural schema-circuits that structure sensory-motor experience as well as abstract understanding. These are acquired very early in life. They are simple, directly understood, and structure complex experiences and concepts. They constitute the semantics of grammar. Cogs include (a) *Event Structures* (“X-nets” (described earlier)) that compute phases of events (such as inception, ongoing, completion, suspension), viewpoints (zoom-in, zoom-out), goals (their achievement and thwarting) and results, outcomes and rewards; (b) *Spatial relations* (schematic image structures (such as containers, orientation, topological relations (inside, outside)); paths (reified trajectories (such as linear, circular)); (c) *Force dynamic interactions* between entities (pushing, pulling holding, releasing, blocking, supporting, helping, hindering, preventing, enabling); (d) *Emotional pathways* (positive and negative), basic emotions (fear, disgust, anger, sadness, happiness, awe, satisfaction, surprise); conceptual schemas for emotions); (e) *Basic entity types and properties* (people, animals, plants, things; substance type; natural vs. artificial; functions; histories); (f) *Basic social relations*; (g) *Quantification* (individuals, pluralities, groups; count vs. mass). [12] has a more detailed discussion of Cogs and their use in language. The computational modeling of Cogs is an ongoing effort within the NTL group at Berkeley. [5,7,12,14] describe the current status of these efforts.

Surface form and Pragmatics

Linguists who have studied narrative have tended to concentrate on the linguistic form of spoken stories, and how narrative structure can be found in them. The modern tradition began with William Labov and Joshua Waletzky’s 1967 paper, “Narrative Analysis: Oral Versions of Personal Experience”[8] and is an excellent example of detailed analysis of surface form in the service of teasing out narrative structure. Figure 5 depicts a recent

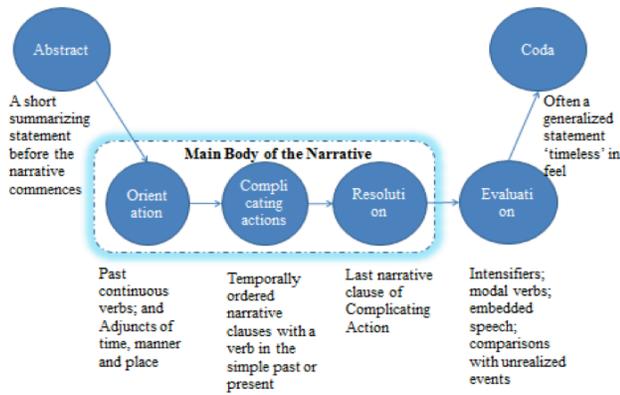


Figure 5: Mapping between linguistic form and structure [8]

version of William Labov’s understanding of how the linguistic form of elementary narratives fit their overall structure:

Figure 5 shows the basic mapping between linguistic form and narrative and plot structure from the work of Labov and collaborators [8]. Other linguistic analyses stress such “pragmatic” factors as viewpoint, presuppositions and implicatures, mental space structure, the subtle meanings of words and grammatical constructions, and so on. Among the prominent authors in this tradition are W. Labov, Charlotte Linde, Alton Becker, Livia Polanyi, Robin Lakoff, and Deborah Tannen (detailed linguistic discussion of these pragmatic factors is outside the scope of this paper).

Conclusion

Narratives exploit the shared cognitive structures of human motivations, goals, emotions, actions, events, and outcomes. Computational models of narrative must therefore be capable of modeling these shared human understandings. We described an ontological framework and an implemented system, KARMA, that is motivated by results in neural computation and captures the fine-grained, context sensitive information about human goals, actions, policies, and outcomes. While the implemented system was able to reason with the richness of the linguistic descriptions in narrative, it did not pay sufficient attention to the inherent form, structure and compositional features of narrative. We identified problems with the existing pilot system and describe several crucial dimensions of narrative structure that have to be modeled in moving toward a computational system that is able to capture the structure and content in human narrative. These additional requirements form the nucleus of ongoing conceptual and modeling work in our laboratory.

References

1. Bryant, John Edward, 2008. Best-Fit Constructional Analysis. Computer Science Division, University of California at Berkeley dissertation.
2. Barrett, Leon 2010. An Architecture for Structured, Concurrent, Real-time Action. Computer Science Division, University of California at Berkeley dissertation.
3. Chang, Nancy, 2009. Constructing grammar: A computational model of the emergence of early constructions. Computer Science Division, University of California at Berkeley dissertation.
4. Chang, Nancy, Srini Narayanan, & Miriam Petruck. 2002a. Putting frames in perspective. In Proc. Nineteenth International Conference on Computational Linguistics (COLING 2002).
5. Feldman, Jerome. 2006. From Molecule to Metaphor. Cambridge, MA: MIT Press.
6. Feldman, Jerome, Ellen Dodge, & John Bryant. 2009. A neural theory of language and embodied construction grammar. In Oxford Handbook of Linguistic Analysis. Oxford University Press.
7. Gallese, Vittorio, & George Lakoff. 2005. The brain’s concepts: the role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology* 22.455–479.
8. Labov W., and J. Waletzky “Narrative Analysis: Oral Versions of Personal Experience” (1968) Special Volume A Journal of Narrative and Life History, Volume 7, 1997, 3-38. <http://www.clarku.edu/~mbamberg/LabovWaletzky.htm>
9. Lakoff, George. 1964 “Structural Complexity in Fairy Tales.” Presented at the Summer Meeting of the Linguistic Society of America. 1964.
10. Lakoff, George. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press.
11. Lakoff, George. 1994. What is metaphor. In *Advances in Connectionist Theory, V3: Analogical Connections*, ed. By Barnden J. & K Holyoak. Addison-Wesley.
12. Lakoff, George. 2009. *The Neural Theory of Metaphor*. Handbook of Metaphor, Cambridge Press, also available at the SSRN eLibrary .
13. Narayanan, Srini, 1997. Knowledge-based Action Representations for Metaphor and Aspect (KARMA). Computer Science Division, University of California at Berkeley dissertation.
14. Narayanan, Srini. 2010. Mind changes: A simulation semantics account of counterfactuals. <http://www.icsi.berkeley.edu/~snarayan/counterfactuals.pdf> (submitted).
15. Propp, Vladimir. *Morphology of the Folktale*. Leningrad, 1928; The Hague: Mouton, 1958; Austin: U. of Texas Press, 1968.
16. Rumelhart, D. E. (1975) Notes on a schema for stories. In D. G. Bobrow and A. Collins (eds) *Representation and Understanding: Studies in Cognitive Science* (New York: Academic Press).
17. Rumelhart, D. E. (1980) On evaluating story grammars. *Cognitive Science*. 4: 313-316.
18. Schank, Roger. C. and Robert P. Abelson, *Scripts, plans, goals and understanding: an inquiry into human knowledge structures*, Erlbaum, 1977.
19. Sinha, Steve, 2008. Answering Questions about Complex Events. Computer Science Division, University of California at Berkeley dissertation.